

Admissibility of solution estimators for stochastic optimization

Amitabh Basu* Tu Nguyen* Ao Sun*

October 21, 2020

Abstract

We look at stochastic optimization problems through the lens of statistical decision theory. In particular, we address admissibility, in the statistical decision theory sense, of the natural sample average estimator for a stochastic optimization problem (which is also known as the *empirical risk minimization (ERM) rule* in learning literature). It is well known that for some simple stochastic optimization problems, the sample average estimator may not be admissible. This is known as *Stein's paradox* in the statistics literature. We show in this paper that for optimizing stochastic linear functions over compact sets, the sample average estimator *is* admissible. Moreover, we study problems with convex quadratic objectives subject to box constraints. Stein's paradox holds when there are no constraints and the dimension of the problem is at least three. We show that in the presence of box constraints, admissibility is recovered for dimensions 3 and 4.

1 Introduction

A large class of stochastic optimization problems can be formulated in the following way:

$$\min_{x \in X} \{f(x) := \mathbb{E}_\xi[F(x, \xi)]\}, \quad (1.1)$$

where $X \subseteq \mathbb{R}^d$ is a fixed feasible region, ξ is a random variable taking values in \mathbb{R}^m , and $F : \mathbb{R}^d \times \mathbb{R}^m \rightarrow \mathbb{R}$. We wish to solve this problem with access to independent samples of ξ . The following are two classical examples:

1. Consider a learning problem with access to labeled samples $(z, y) \in \mathbb{R}^n \times \mathbb{R}$ from some distribution and the goal is to find a function $f \in \mathcal{F}$ in a finitely parametrized hypothesis class \mathcal{F} (e.g., all neural network functions with a fixed architecture) that minimizes expected loss, where the loss function is given by $\ell : \mathbb{R} \times \mathbb{R} \rightarrow \mathbb{R}_+$. One can model this using (1.1) by setting d to be the number of parameters for \mathcal{F} , $m = n + 1$, $X \subseteq \mathbb{R}^d$ is the subset that describes \mathcal{F} via the parameters, and $F(f, (z, y)) = \ell(f(z), y)$.

*Department of Applied Mathematics and Statistics, The Johns Hopkins University. Amitabh Basu and Tu Nguyen were supported by NSF grant CMMI1452820 and ONR grant N000141812096. Amitabh Basu also acknowledges support from NSF grant CCF1934979.

2. When $d = m$, $F(x, \xi) = \|x - \xi\|^2$, $X = \mathbb{R}^d$ and ξ is distributed with mean μ , (1.1) becomes

$$\min_{x \in \mathbb{R}^d} \mathbb{E}_\xi[\|x - \xi\|^2] = \min_{x \in \mathbb{R}^d} \|x - \mathbb{E}[\xi]\|^2 + \mathbb{V}[\xi]$$

In particular, if one knows $\mu := \mathbb{E}[\xi]$, the optimal solution is given by $x = \mu$. Thus, this stochastic optimization problem becomes equivalent to the classical statistics problem of estimating the mean of the distribution of ξ , given access to independent samples.

We would like to emphasize our data-driven viewpoint on the problem (1.1). In particular, we will not assume detailed knowledge of the distribution of the random variable ξ , but only assume that it comes from a large family of distributions. More specifically, we will not assume knowledge of means or higher order moments, and certainly not the exact distribution of ξ . This is in contrast to some approaches within the stochastic optimization literature that proceed on the assumption that such detailed knowledge of the distribution is at hand. Such an approach would rewrite (1.1) by finding an analytic expression for the expectation $\mathbb{E}_\xi[F(x, \xi)]$ (in terms of the known parameters of the distribution of ξ), perhaps with some guaranteed approximation if an exact analysis is difficult. The problem then becomes a *deterministic* optimization problem¹, often a very complicated and difficult one, which is then attacked using novel and innovative ideas of mathematical optimization. See [5] for a textbook exposition of this viewpoint.

In contrast, as mentioned above, we will assume that the true distribution of ξ comes from a postulated large family of (structured) distributions, and we assume that we have access to data points ξ^1, ξ^2, \dots drawn independently from the true distribution of ξ . This makes our approach distinctly statistical and data-driven in nature. We “learn” or glean information about the distribution of ξ from the data, which we then use to “solve” (1.1). Statistical decision theory becomes a natural framework for such a viewpoint, to formalize what it even means to “solve” the problem after “learning” about the distribution from data. We briefly review relevant terminology from statistical decision theory below.

We do not mean to imply that a statistical perspective on stochastic optimization has never been studied before this paper. This is far from true; see [5, Chapter 9] and [27, Chapter 5] for detailed discussions of statistical approaches and methods in stochastic optimization. In [25] and [8], the authors introduce a statistical decision theory perspective that is essentially the same as our framework. In recent parlance, “data-driven optimization” has been used to describe the statistical viewpoint and has a vast literature; some recent papers closely related to our work are [2, 11, 12, 17, 30], with [11, 17] particularly close in spirit to this paper. Nevertheless, our perspective is quite different and follows in the footsteps of the inspirational paper of Davarnia and Cornuejols [9].

1.1 Statistical decision theory and admissibility

Statistical decision theory is a mathematical framework for modeling decision making in the face of uncertain or incomplete information. One models the uncertainty by a set of *states*

¹Note that we are *not* referring to what stochastic optimization literature refers to as the deterministic model where all appearances of any random variable in the problem are replaced by its expectation. We are really talking about what the stochastic optimization community would refer to as the stochastic problem and a stochastic solution.

of nature denoted by Θ . The decision making process is to choose an *action* from a set \mathcal{A} that performs best in a given state of nature θ . To take our stochastic optimization setting, the set of states of nature is given by the family \mathcal{D} of distributions that we believe the true distribution of ξ comes from, and the set of actions \mathcal{A} is the feasible region X , i.e., select $x \in X$ that minimizes $f(x) := \mathbb{E}_{\xi \sim D}[F(x, \xi)]$. In the general framework of decision theory, one defines a *loss function* $\mathcal{L} : \Theta \times \mathcal{A} \rightarrow \mathbb{R}_+$ to evaluate the performance of an action $a \in \mathcal{A}$ against a state of nature $\theta \in \Theta$. The smaller $\mathcal{L}(\theta, a)$ is, the better $a \in \mathcal{A}$ does with respect to the state $\theta \in \Theta$ ². In our setting of stochastic optimization, we take an action $\hat{x} \in X$. The natural way to evaluate its performance is via the so-called *optimality gap*, i.e., how close is $f(\hat{x})$ to the optimal value of (1.1). Therefore, the following is a natural loss function for stochastic optimization:

$$\begin{aligned}\mathcal{L}(D, x) &:= f(x) - f(x(D)) \\ &= \mathbb{E}_{\xi \sim D}[F(x, \xi)] - \mathbb{E}_{\xi \sim D}[F(x(D), \xi)],\end{aligned}\tag{1.2}$$

where $x(D)$ is an optimal solution to (1.1) when $\xi \sim D$.

The *statistical* aspect of statistical decision theory comes from the fact that the state θ is not revealed directly, but only through data/observations based on θ that can be noisy or incomplete. This is formalized by postulating a parameterized family of probability distributions $\mathcal{P} := \{P_\theta : \theta \in \Theta\}$ on a common *sample space* \mathcal{X} . After observing a realization $y \in \mathcal{X}$ of this random variable, one forms an opinion about what the possible state is and one chooses an action $a \in \mathcal{A}$. Formally, a *decision rule* is a function $\delta : \mathcal{X} \rightarrow \mathcal{A}$ giving an action $\delta(y) \in \mathcal{A}$ when data $y \in \mathcal{X}$ is observed. To take our particular setting of stochastic optimization, one observes data points $\xi^1, \xi^2, \dots, \xi^n$ that are i.i.d. realizations of $\xi \sim D$; thus, $\mathcal{X} = \underbrace{\mathbb{R}^m \times \mathbb{R}^m \times \dots \times \mathbb{R}^m}_{n \text{ times}}$ with distributions $\mathcal{P} := \underbrace{\{D \times D \times \dots \times D : D \in \mathcal{D}\}}_{n \text{ times}}$ on \mathcal{X} parameterized by the states $D \in \mathcal{D}$.

Finally, one evaluates decision rules by averaging over the data, defining the *risk function*

$$\mathcal{R}(\theta, \delta) := \mathbb{E}_{y \sim P_\theta}[\mathcal{L}(\theta, \delta(y))].$$

One can think of the risk function as mapping a decision rule to a nonnegative function on the class of distributions \mathcal{P} ; this function is sometimes called the *risk of the decision rule*. A decision rule is “good” if its risk has “low” values. A basic criterion for choosing decision rules is then the following. We say that δ' *weakly dominates* δ if $\mathcal{R}(\theta, \delta') \leq \mathcal{R}(\theta, \delta)$ for all $\theta \in \Theta$. We say that δ' *dominates* δ if, in addition, $\mathcal{R}(\hat{\theta}, \delta') < \mathcal{R}(\hat{\theta}, \delta)$ for some $\hat{\theta} \in \Theta$. A decision rule δ is said to be *inadmissible* if there exists another decision rule δ' that dominates δ . A decision rule δ is said to be *admissible* if it is not dominated by any other decision rule. In-depth discussions of general statistical decision theory can be found in [1, 4, 24].

²We caution the reader that the use of the words “loss” and “risk” in statistical decision theory are somewhat different from their use in machine learning literature. In machine learning, the function $F(x, \xi)$ is usually referred to as “loss” and the function $f(x)$ is referred to as “risk” in (1.1). Thus Example 1. above becomes a “risk minimization” problem with an associated “empirical risk minimization (ERM)” problem when one replaces the expectation by a sample average.

1.2 Admissibility of the sample average estimator and our results

We would like to study the admissibility of natural decision rules for solving (1.1). As explained above, we put this in the decision theoretical framework by setting the sample space $\mathcal{X} = \underbrace{\mathbb{R}^m \times \dots \times \mathbb{R}^m}_{n \text{ times}}$, where n is the number of i.i.d. observations one makes for $\xi \sim D$ for $D \in \mathcal{D}$, and \mathcal{D} is a fixed family of distributions. A decision rule is now a map $\delta : \underbrace{\mathbb{R}^m \times \mathbb{R}^m \times \dots \times \mathbb{R}^m}_{n \text{ times}} \rightarrow X$. The class of distributions on \mathcal{X} is $\mathcal{P} = \{\underbrace{D \times D \times \dots \times D}_{n \text{ times}} : D \in \mathcal{D}\}$. The loss function is defined as in (1.2). In this paper, we wish to study the admissibility of the *sample average* decision rule δ_{SA} defined as

$$\delta_{SA}(\xi^1, \dots, \xi^n) \in \arg \min \left\{ \frac{1}{n} \sum_{i=1}^n F(x, \xi^i) : x \in X \right\} \quad (1.3)$$

This is a standard procedure in stochastic optimization, and often goes by the name of *sample average approximation (SAA)*; in machine learning, it goes by the name of *Empirical Risk Minimization (ERM)*. To emphasize the dependence on the number of samples n , we introduce a superscript, i.e., δ_{SA}^n will denote the estimator based on the sample average of the objective from n observations. Moreover, for any $n \in \mathbb{N}$, let Δ^n be the set of all decision rules $\delta : \underbrace{\mathbb{R}^m \times \mathbb{R}^m \times \dots \times \mathbb{R}^m}_{n \text{ times}} \rightarrow X$ such that $\mathbb{E}_{\xi^1, \dots, \xi^n}[\delta(\xi^1, \dots, \xi^n)]$ exists.

Stein's paradox. It turns out that there are simple instances of problem (1.1) where the sample average rule is *inadmissible*. Consider the setting of Example 2 in the Introduction, where $m = d$, $F(x, \xi) = \|x - \xi\|^2$ and $X = \mathbb{R}^d$. We assume ξ is distributed normally $\xi \sim N(\mu, I)$ with unknown mean μ ; here I denotes the identity matrix. In the language of statistical decision theory, the states of nature are now parametrized by $\mu \in \mathbb{R}^d$. Since $\mathbb{E}_{\xi \sim D}[F(x, \xi)] = \|x - \mathbb{E}[\xi]\|^2 + \mathbb{V}[\xi]$, the minimizer is simply $x(D) = \mathbb{E}[\xi] = \mu$ with objective value $\mathbb{V}[\xi]$. Evaluating (1.2), $\mathcal{L}(\mu, x) = \mathbb{E}_{\xi \sim D}[F(x, \xi)] - \mathbb{E}_{\xi \sim D}[F(x(D), \xi)] = (\|x - \mathbb{E}[\xi]\|^2 + \mathbb{V}[\xi]) - \mathbb{V}[\xi] = \|x - \mu\|^2$. Thus, $x = \mu$ minimizes the loss when the state of nature is μ . Consequently, the problem becomes the classical problem of estimating the mean of a Gaussian from samples under ‘‘squared distance loss’’. Also, the sample average decision rule solves (1.3) which is the problem $\min\{\frac{1}{n} \sum_{i=1}^n \|x - \xi^i\|^2 : x \in \mathbb{R}^d\}$ and therefore returns the empirical average of the samples, i.e., $\delta_{SA}^n(\xi^1, \dots, \xi^n) = \bar{\xi}$ where $\bar{\xi} := \frac{1}{n} \sum_{i=1}^n \xi^i$. It is well-known that this sample average decision rule is *inadmissible* if $d \geq 3$; this was first observed by Stein [28] and is commonly referred to as *Stein's paradox* in statistics literature. The *James-Stein estimator* [20] strictly dominates the sample average estimator; see [1, 24] for an exposition.

Our results. We focus on two particular cases of the stochastic optimization problem (1.1):

1. $m = d$, $F(x, \xi) = \xi^T x$, $X \subseteq \mathbb{R}^d$ is a given compact (not necessarily convex) set, and ξ has a *Gaussian distribution* with *unknown* mean $\mu \in \mathbb{R}^d$ and covariance $\Sigma \in \mathbb{R}^{d \times d}$ denoted by $\xi \sim N(\mu, \Sigma)$. In other words, we optimize an uncertain linear objective

over a fixed compact set. Note that, along with linear or convex optimization, we also capture non-convex feasible regions like mixed-integer non-linear optimization or linear complementarity constraints.

2. $m = d$, $F(x, \xi) = \frac{1}{2}\|x\|^2 - \xi^T x$, $X \subseteq \mathbb{R}^d$ is a box constrained set, i.e., $X := \{x \in \mathbb{R}^d : \ell_i \leq x_i \leq u_i, i = 1, \dots, d\}$ ($\ell_i \leq u_i$ are arbitrary real numbers), and ξ has a *Gaussian distribution* with *unknown* mean $\mu \in \mathbb{R}^d$ and covariance $\Sigma \in \mathbb{R}^{d \times d}$ denoted by $\xi \sim N(\mu, \Sigma)$. Here, we wish to minimize a convex quadratic function with an uncertain linear term over box constraints.

In the first case, we show that there is no “Stein’s paradox” type phenomenon, i.e., the sample average solution is admissible for every $d \in \mathbb{N}$. For the second case, we show that the sample average solution is admissible for $d \leq 4$. Note that in the second situation above, $F(x, \xi) = \frac{1}{2}\|x - \xi\|^2 - \frac{1}{2}\|\xi\|^2$ and thus the problem of minimizing $\mathbb{E}[\frac{1}{2}\|x - \xi\|^2 + \frac{1}{2}\|\xi\|^2] = \mathbb{E}[\frac{1}{2}\|x - \xi\|^2] + \frac{1}{2}\mathbb{E}[\|\xi\|^2]$ is equivalent to the setting of Stein’s paradox (since $\frac{1}{2}\mathbb{E}[\|\xi\|^2]$ is just a constant), except that we now impose box constraints on x . Thus, admissibility is recovered for $d = 3, 4$ with box constraints. While we are unable to establish it for $d \geq 5$, we strongly suspect that there is no Stein’s paradox in *any dimension* once box constraints are imposed; in fact, we believe this is true when *any* compact constraint set is imposed (see discussion below). The precise statements of our results follow.

THEOREM 1.1. *Consider problem (1.1) in the setting where X is a given compact set and $F(\xi, x) = \xi^T x$, and $\xi \sim N(\mu, \Sigma)$ with unknown μ and Σ . The sample average rule now simply becomes*

$$\delta_{SA}^n(\xi^1, \dots, \xi^n) \in \arg \min \{\bar{\xi}^T x : x \in X\} \quad (1.4)$$

where $\bar{\xi} := \frac{1}{n} \sum_{i=1}^n \xi^i$ denotes the sample average of the observed objective vectors. For any $n \in \mathbb{N}$, and any $\Sigma \in \mathbb{R}^{d \times d}$, we consider the states of nature to be parametrized by $\mu \in \mathbb{R}^d$. Then for every $n \in \mathbb{N}$ and $\Sigma \in \mathbb{R}^{d \times d}$, δ_{SA}^n is admissible within Δ^n .

THEOREM 1.2. *Let $d \leq 4$. Consider problem (1.1) in the setting where $X := \{x \in \mathbb{R}^d : \ell_i \leq x_i \leq u_i, i = 1, \dots, d\}$ ($\ell_i \leq u_i$ are arbitrary real numbers) and $F(\xi, x) = \frac{1}{2}\|x\|^2 - \xi^T x$, and $\xi \sim N(\mu, \Sigma)$ with unknown μ and Σ . The sample average rule now simply becomes*

$$\delta_{SA}^n(\xi^1, \dots, \xi^n) \in \arg \min \left\{ \frac{1}{2}\|x\|^2 - \bar{\xi}^T x : x \in X \right\} \quad (1.5)$$

where $\bar{\xi} := \frac{1}{n} \sum_{i=1}^n \xi^i$ denotes the sample average of the observed vectors. For any $n \in \mathbb{N}$, and any $\Sigma \in \mathbb{R}^{d \times d}$, we consider the states of nature to be parametrized by $\mu \in \mathbb{R}^d$. Then for every $n \in \mathbb{N}$ and $\Sigma \in \mathbb{R}^{d \times d}$, δ_{SA}^n is admissible within Δ^n .

The following result provides some concrete basis for our belief that compact constraints imply admissibility in any dimension even in the quadratic case; alas, we are unable to establish it even for box constraints when $d \geq 5$.

THEOREM 1.3. *When $X = \{x \in \mathbb{R}^d : \|x\| \leq R\}$ for some $R > 0$, $F(\xi, x) = \frac{1}{2}\|x\|^2 - \xi^T x$, and $\xi \sim N(\mu, \Sigma)$, then the sample average rule is admissible for all $d \geq 1$.*

We present two different proofs of Theorem 1.1. The first one, presented in Sections 3 and 3.2 uses a novel proof technique for admissibility, to the best of our knowledge. The second proof, presented in Section 4 uses the conventional idea of showing that the sample average estimator is the (unique) Bayes estimator under an appropriate prior. We feel that the first proof technique could be useful for future research into the question of admissibility of solution estimators for stochastic optimization. The second method using Bayes estimators is easier to generalize to the quadratic settings of Theorems 1.2 and 1.3, and thus forms a natural segue into their proofs presented in Section 5.

1.3 Comparison with previous work

The statistical decision theory perspective on stochastic optimization presented here follows the framework of [9] and [10]. In particular, the authors of [9] consider admissibility of solution estimators in two different stochastic optimization problems: one where $X = \mathbb{R}^d$ and $F(x, \xi) = x^T Qx + \xi^T x$ for some fixed matrix positive definite matrix Q (i.e., unconstrained convex quadratic minimization), and the second one where X is the unit ball and $F(x, \xi) = \xi^T x$. ξ is again assumed to be distributed according to a normal distribution $N(\mu, I)$ with unknown mean μ . They show that the sample average approximation is *not* admissible in general for the first problem, and it is admissible for the second problem. Note that the second problem is a special case of our setting. In both these cases, there is a closed-form solution to the deterministic version of the optimization problem, which helps in the analysis. This is not true for the general optimization problem we consider here (even if we restrict X to be a polytope, we get a linear program which, in general, has no closed form solution).

Another difference between our work and [9] is the following. In [9], the question of admissibility is addressed within a smaller subset of decision rules that are “decomposable” in the sense that any decision rule is of the form $\tau \circ \kappa$, where $\kappa : \underbrace{\mathbb{R}^d \times \mathbb{R}^d \times \dots \times \mathbb{R}^d}_{n \text{ times}} \rightarrow \mathbb{R}^d$

maps the data ξ^1, \dots, ξ^n to a vector $\hat{\mu} \in \mathbb{R}^d$ and then $\tau : \mathbb{R}^d \rightarrow X$ is of the form $\tau(\hat{\mu}) \in \arg \min \{\hat{\mu}^T x : x \in X\}$. In other words, one first estimates the mean of the uncertain objective (using any appropriate decision rule) and then uses this estimate to solve a deterministic optimization problem. In the follow-up work [10], the authors call such decision rules *Separate estimation-optimization (Separate EO) schemes* and more general decision rules as *Joint estimation-optimization (Joint EO) schemes*. In this paper, we establish admissibility of the sample average estimator within general decision rules (joint EO schemes in the terminology of [10]). The only condition we put on the decision rules is that of integrability, which is a minimum requirement needed to even define the risk of a decision rule. Note that proving inadmissibility within separate EO schemes implies inadmissibility within joint EO schemes. On the other hand, establishing admissibility within joint EO schemes means defending against a larger class of decision rules. The general concept of joint estimation-optimization schemes also appears in [8, 11, 25], presented in slightly different vocabulary.

As mentioned before, the quadratic convex objective has been studied in statistics in the large body of work surrounding Stein’s paradox, albeit not in the stochastic optimization language that we focus on here. Moreover, all of this classical work is for the unconstrained problem. To the best of our knowledge, the version with box constraints has not been studied before (but see [3, 6, 7, 13, 15, 16, 18, 22, 23, 26] and the book [14] for a related, but different,

statistical problem that has received a lot of attention). It is very intriguing (at least to us) that in the presence of such constraints, admissibility is recovered for dimensions $d = 3, 4$; recall that for the unconstrained problem, the sample average solution is admissible only for $d = 1, 2$ and inadmissible for $d \geq 3$.

1.4 Admissibility and other notions of optimality

We end our discussion of the results with a few comments about other optimality notions for decision rules. In large sample statistics, one often considers the behavior of decision rules when $n \rightarrow \infty$ (recall n is the number of samples). A sequence of decision rules δ^n (each δ^n is based on n i.i.d samples) is said to be *asymptotically inadmissible* if there exists a decision rule sequence $\bar{\delta}^n$ such that $\lim_{n \rightarrow \infty} \frac{R(\theta, \bar{\delta}^n)}{R(\theta, \delta^n)} \leq 1$ for every θ , and for some $\hat{\theta}$ the limiting ratio is strictly less than 1. Admissibility for every $n \in \mathbb{N}$ does not necessarily imply asymptotic admissibility [4, Problem 4, page 200], and asymptotic admissibility does not imply admissibility for finite $n \in \mathbb{N}$, i.e., there can be decision rules that are inadmissible for every $n \in \mathbb{N}$ and yet be asymptotically admissible. Thus, the small-sample behaviour (fixed $n \in \mathbb{N}$) and large sample behavior ($n \rightarrow \infty$) can be quite different. One advantage of proving asymptotic admissibility is that it also implies the *rate* of convergence of the risk (as a function of n) is optimal; such rules are called *rate optimal*. Unfortunately, our admissibility results about δ_{SA}^n do not immediately imply asymptotic admissibility or rate optimality. Standard techniques for proving rate optimality such as Hajek-Le Cam theory [29, Chapter 8] cannot be applied because regularity assumptions about the decision rules are not satisfied in our setting. For example, in the linear objective case discussed above, δ_{SA}^n has a degenerate distribution that is supported on the boundary of the feasible region X . Similarly, in the quadratic case, δ_{SA}^n 's distribution is supported on the compact domain X , with majority of the mass on the boundary when μ is outside X . This rules out any possibility of “asymptotic normality” or “local asymptotic minimaxity” results [29, Chapters 7, 8].

While we are unable to prove rate optimality for δ_{SA}^n , it is reasonably straightforward to show that δ_{SA}^n is *consistent* in the sense that $R(\theta, \delta_{SA}^n) \rightarrow 0$ as $n \rightarrow \infty$. This can be derived from consistency results in stochastic optimization literature [27, Chapter 5], but we present the argument here for completeness. In the linear objective case, the loss for δ_{SA}^n is given by $\mathcal{L}(\mu, \delta_{SA}^n) = \mu^T \delta_{SA}^n - \mu^T x(\mu) = (\mu - \bar{\xi})^T \delta_{SA}^n + \bar{\xi}^T \delta_{SA}^n - \mu^T x(\mu) \leq (\mu - \bar{\xi})^T \delta_{SA}^n + \bar{\xi}^T x(\mu) - \mu^T x(\mu)$ since δ_{SA}^n is the minimizer with respect to $\bar{\xi}$. Thus, $\mathcal{L}(\mu, \delta_{SA}^n) \leq (\mu - \bar{\xi})^T (\delta_{SA}^n - x(\mu)) \leq \|\mu - \bar{\xi}\| \cdot K$ by the Cauchy-Schwarz inequality, where K is the diameter of the compact feasible region X . Therefore, $R(\mu, \delta_{SA}^n) \leq K \mathbb{E}[\|\mu - \bar{\xi}\|]$. Since the sample average $\bar{\xi}$ has a normal distribution with mean μ and variance that scales like $O(1/n)$, $R(\mu, \delta_{SA}^n) \rightarrow 0$ with rate $O(1/\sqrt{n})$. A similar argument can be made in the quadratic objective case. However, we are unable to show that $O(1/\sqrt{n})$ is the optimal rate in either case.

There is a large body of literature on *shrinkage estimators* in the unconstrained, quadratic objective setting. A relatively recent insight [31] shows that as $d \rightarrow \infty$ (recall d is the dimension), a certain class of shrinkage estimators (called *SURE estimators*) have risk functions that dominate any other shrinkage estimator's risk, and hence the sample average estimator's risk, with just a single sample. This potentially suggests that the phenomenon presented here, where admissibility of δ_{SA}^n is recovered for $d = 3, 4$, holds only for small dimensions and for large enough dimensions, the sample average estimator remains inadmissible. However, this

is not immediate because of two reasons: 1) the value of d for which the SURE estimator in [31] starts to dominate any other estimator depends on the parameter μ , and 2) the setting in [31] is still unconstrained optimization. In fact, as stated earlier, we strongly suspect that with compact constraints, admissibility of δ_{SA}^n holds for *all* dimensions (see Theorem 1.3). We remark that SURE estimators need not be admissible themselves [21].

There are other notions of optimality of decision rules even in the small/finite sample setting. For example, the *minimax* decision rule minimizes the sup norm of the risk function, i.e., one solves $\inf_{\delta} \sup_{\theta} R(\theta, \delta)$. In general, admissibility does not imply minimaxity, nor does minimaxity imply admissibility. Of course, if a minimax rule is inadmissible, then the dominating rule is also minimax and is certainly to be preferred, unless computational concerns prohibit this. In many settings however (e.g., estimation in certain exponential and group families [24, Chapter 5]), minimax rules are also provably admissible and thus minimaxity is a more desirable criterion.

Generally speaking, admissibility is considered a weak notion of optimality because admissible rules may have undesirable properties like very high risk values for certain states of the world. Moreover, as noted above, inadmissible rules may have optimal large sample behavior. Nevertheless, it is useful to know if widely used decision rules such as sample average approximations satisfy the basic admissibility criterion, because if not, then one could use the dominating decision rule unless it is computationally much more expensive.

2 Technical Tools

We first recall a basic fact from calculus.

LEMMA 2.1. *Let $F : \mathbb{R}^m \rightarrow \mathbb{R}$ be a twice continuously differentiable map such that $F(0) = 0$. Suppose $\nabla^2 F(0)$ is not negative semidefinite; in other words, there is a direction $d \in \mathbb{R}^d$ of positive curvature, i.e., $d^T \nabla^2 F(0) d > 0$. Then there exists $z \in \mathbb{R}^m$ such that $F(z) > 0$.*

Proof. If $\nabla F(0) \neq 0$, then there exist $\lambda > 0$ such that $F(z) > 0$ for $z = \lambda \nabla F(0)$ since $F(0) = 0$. Else, if $\nabla F(0) = 0$ then there exists $\lambda > 0$ such that $F(z) > 0$ for $z = \lambda d$, where d is the direction of positive curvature at 0. \square

We will need the following central definition and result from statistics. See e.g., Section 6, Chapter 1 in [24].

DEFINITION 2.2. *A statistic is a function $T : \chi \rightarrow \mathbb{R}^m$, i.e., it is any function that maps the data to a vector (or a scalar if $m = 1$). Let \mathcal{P} be a family of distributions on the sample space χ . A sufficient statistic for \mathcal{P} is a statistic on χ such that the conditional distribution on χ given $T = t$ does not depend on the distribution from \mathcal{P} , for all $t \in \mathbb{R}^m$.*

PROPOSITION 2.3. *Let $\chi = \underbrace{\mathbb{R}^d \times \mathbb{R}^d \times \dots \times \mathbb{R}^d}_{n \text{ times}}$ and let $\mathcal{P} = \underbrace{\{N(\mu, I) \times N(\mu, I) \times \dots \times N(\mu, I)\}}_{n \text{ times}}$: $\mu \in \mathbb{R}^d\}$, i.e., $(\xi^1, \dots, \xi^n) \in \chi$ are i.i.d samples from the normal distribution $N(\mu, I)$. Then $T(\xi^1, \dots, \xi^n) = \bar{\xi} := \frac{1}{n} \sum_{i=1}^n \xi^i$ is a sufficient statistic for \mathcal{P} .*

We will also need the following useful property for the family of normal distributions $\{N(\mu, I) : \mu \in \mathbb{R}^m\}$. Indeed the following result is true for any *exponential family* of distributions; see Theorem 5.8, Chapter 1 in [24] for details.

THEOREM 2.4. *Let $f : \mathbb{R}^m \rightarrow \mathbb{R}^d$ be any integrable function. The function*

$$h(\mu) := \int_{\mathbb{R}^m} f(y) e^{-\frac{n\|y-\mu\|^2}{2}} dy$$

is continuous and has derivatives of all orders with respect to μ , which can be obtained by differentiating under the integral sign.

In the rest of the paper, for any vector v , v_j will denote the j -th coordinate, and for any matrix $A \in \mathbb{R}^{p \times q}$, A_{ij} will denote entry in the i -th row and j -th column.

We need one further result on the geometry of the hypercube which is easy to verify. We recall that for any closed, convex set $C \subseteq \mathbb{R}^d$ and point $x \in C$, the *normal cone at x* is defined to be the set of all vectors $c \in \mathbb{R}^d$ such that $x \in \operatorname{argmax}_{y \in C} c^T y$. We extend this concept to any face F of C : The *normal cone at the face F* is defined to be the set of all vectors $c \in \mathbb{R}^d$ such that $F \subseteq \operatorname{argmax}_{y \in C} c^T y$.

LEMMA 2.5. *Let $X = \{x \in \mathbb{R}^d : -u_i \leq x_i \leq u_i \quad i = 1, \dots, d\}$ be a box centered at the origin. For any face $F \subseteq X$ of X (possibly with $F = X$), let $I_F^+ \subseteq \{1, \dots, d\}$ be the subset of coordinates which are set to the bound u_i for all points in F , $I_F^- \subseteq \{1, \dots, d\}$ be the subset of coordinates which are set to the bound $-u_i$ for all points in F , and N_F denote the normal cone at F . Then the following are true:*

1. *For any face $F \subseteq X$,*

$$F + N_F = \left\{ x \in \mathbb{R}^d : \begin{array}{ll} x_i \geq u_i & i \in I_F^+ \\ x_i \leq -u_i & i \in I_F^- \\ -u_i \leq x_i \leq u_i & i \notin I_F^+ \cup I_F^- \end{array} \right\}.$$

2. *The interior of $F + N_F$ is disjoint from the interior of $F' + N_{F'}'$ whenever $F \neq F'$ and we have the following decomposition of \mathbb{R}^d :*

$$\mathbb{R}^d = \bigcup_{F \text{ face of } X} F + N_F$$

3 Proof of Theorem 1.1 (the scenario with linear objective)

3.1 When the covariance matrix is the identity

Proof of Theorem 1.1 when $\Sigma = I$. As introduced in the previous sections, $\bar{\xi}$ will denote the sample average of ξ^1, \dots, ξ^n . Consider an arbitrary decision rule $\delta \in \Delta^n$. Consider the conditional expectation

$$\eta(y) = \mathbb{E}_{\xi^1, \dots, \xi^n} [\delta(\xi^1, \dots, \xi^n) | \bar{\xi} = y].$$

Observe that $\eta(y) \in \operatorname{conv}(X)$ (i.e., the convex hull of X , which is compact since X is compact) since δ maps into X . Moreover, since $\bar{\xi}$ is a sufficient statistic for the family of normal distributions by Proposition 2.3, $\eta(y)$ does not depend on μ . This is going to be important below. To maintain intuitive notation, we will also say that δ_{SA}^n is given by $\delta_{SA}^n(\xi^1, \dots, \xi^n) =$

$\eta^*(\bar{\xi})$, where $\eta^*(y)$ returns a point in $\arg \min \{y^T x : x \in X\}$. Note also that for any action $x \in X$, (1.2) evaluates to

$$\mathcal{L}(\mu, x) = \mu^T x - \mu^T x(\mu),$$

where $x(\mu)$ denotes the optimal solution to the problem $\min \{\mu^T x : x \in X\}$. Using the law of total expectation,

$$\begin{aligned} R(\mu, \delta) &= \mathbb{E}_{\xi^1, \dots, \xi^n} [\mathcal{L}(\mu, \delta(\xi^1, \dots, \xi^n))] \\ &= \mathbb{E}_{\xi^1, \dots, \xi^n} [\mu^T \delta(\xi^1, \dots, \xi^n)] - \mu^T x(\mu) \\ &= \mathbb{E}_y [\mathbb{E}_{\xi^1, \dots, \xi^n} [\mu^T \delta(\xi^1, \dots, \xi^n) | \bar{\xi} = y]] - \mu^T x(\mu) \\ &= \mathbb{E}_y [\mu^T \eta(y)] - \mu^T x(\mu) \end{aligned}$$

If $\eta = \eta^*$ almost everywhere, then $R(\mu, \delta) = R(\mu, \delta_{SA}^n)$ for all $\mu \in \mathbb{R}^d$, and we would be done. So in the following, we assume that $\eta \neq \eta^*$ on a set of strictly positive measure. This implies the following

CLAIM 3.1. *For all $y \in \mathbb{R}^d$, $y^T \eta(y) \geq y^T \eta^*(y)$ and the set $\{y \in \mathbb{R}^d : y^T \eta(y) > y^T \eta^*(y)\}$ is of strictly positive measure.*

Proof. Since X is compact, $\text{conv}(X)$ is a compact, convex set and $\min\{y^T x : x \in \text{conv}(X)\} = \min\{y^T x : x \in X\}$ for every $y \in \mathbb{R}^d$. Therefore, since $\eta(y) \in \text{conv}(X)$ and $\eta^*(y) \in \arg \min\{y^T x : x \in X\}$, we have $y^T \eta(y) \geq y^T \eta^*(y)$ for all $y \in \mathbb{R}^d$.

Since $\text{conv}(X)$ is a compact, convex set, the set of $y \in \mathbb{R}^d$ such that $|\arg \min\{y^T x : x \in \text{conv}(X)\}| > 1$ is of zero Lebesgue measure. Let $S \subseteq \mathbb{R}^d$ be the set of $y \in \mathbb{R}^d$ such that $\arg \min\{y^T x : x \in \text{conv}(X)\}$ is a singleton, i.e., there is a unique optimal solution; so $\mathbb{R}^d \setminus S$ has zero Lebesgue measure. Let $D := \{y \in \mathbb{R}^d : \eta(y) \neq \eta^*(y)\}$. Since we assume that D has strictly positive measure, $D \cap S$ must have strictly positive measure. Consider any $y \in D \cap S$. Since $\min\{y^T x : x \in X\} = \min\{y^T x : x \in \text{conv}(X)\}$, we must have $\arg \min\{y^T x : x \in X\} \subseteq \arg \min\{y^T x : x \in \text{conv}(X)\}$. Since $y \in S$, $\arg \min\{y^T x : x \in \text{conv}(X)\}$ is a singleton and thus $\eta^*(y)$ is the unique optimum for $\min\{y^T x : x \in \text{conv}(X)\}$. Since $y \in D$, $\eta(y) \neq \eta^*(y)$, and therefore $y^T \eta(y) > y^T \eta^*(y)$. Thus, we have the second part of the claim. \square

Now consider the function $F : \mathbb{R}^d \rightarrow \mathbb{R}$ defined by

$$F(\mu) := R(\mu, \delta) - R(\mu, \delta_{SA}^n). \quad (3.1)$$

To show that δ_{SA}^n is admissible, it suffices to show that there exists $\bar{\mu} \in \mathbb{R}^d$ such that $F(\bar{\mu}) > 0$. For any $\mu \in \mathbb{R}^d$, we have from above

$$\begin{aligned} F(\mu) &= R(\mu, \delta) - R(\mu, \delta_{SA}^n) \\ &= \mathbb{E}_y [\mu^T \eta(y)] - \mathbb{E}_y [\mu^T \eta^*(y)] \\ &= \mu^T \int_{\mathbb{R}^d} (\frac{n}{2\pi})^{n/2} \eta(y) e^{\frac{-n\|y-\mu\|^2}{2}} dy - \mu^T \int_{\mathbb{R}^d} (\frac{n}{2\pi})^{n/2} \eta^*(y) e^{\frac{-n\|y-\mu\|^2}{2}} dy \\ &= \mu^T \int_{\mathbb{R}^d} (\frac{n}{2\pi})^{n/2} (\eta(y) - \eta^*(y)) e^{\frac{-n\|y-\mu\|^2}{2}} dy \end{aligned}$$

where in the second to last equality, we have used the fact that $\bar{\xi}$ has distribution $N(\mu, \frac{1}{n}I)$. Note that the formula above immediately gives $F(0) = 0$. We will employ Lemma 2.1 on $F(\mu)$

to show the existence of $\bar{\mu} \in \mathbb{R}^d$ such that $F(\bar{\mu}) > 0$. For this purpose, we need to compute the gradient $\nabla F(\mu)$ and Hessian $\nabla^2 F(\mu)$. We alert the reader that in these calculations, it is crucial that $\eta(y)$ does not depend on μ (due to sufficiency of the sample average) and hence it is to be considered as a constant when computing the derivatives below. For ease of calculation, we introduce the following functions $E, G^1, \dots, G^d : \mathbb{R}^d \rightarrow \mathbb{R}^d$:

$$\begin{aligned} E(\mu) &:= (\frac{n}{2\pi})^{n/2} \int_{\mathbb{R}^d} (\eta(y) - \eta^*(y)) e^{-\frac{n\|y-\mu\|^2}{2}} dy, \\ G^i(\mu) &:= (\frac{n}{2\pi})^{n/2} \int_{\mathbb{R}^d} y_i (\eta(y) - \eta^*(y)) e^{-\frac{n\|y-\mu\|^2}{2}} dy, \\ i &= 1, \dots, d. \end{aligned}$$

So $F(\mu) = \mu^T E(\mu)$. We also define the map $G : \mathbb{R}^d \rightarrow \mathbb{R}^{d \times d}$ as

$$G(\mu)_{ij} = (G^i(\mu))_j.$$

CLAIM 3.2. *For any $\mu \in \mathbb{R}^d$, $\nabla F(\mu) = E(\mu) + nG(\mu)\mu - n(\mu^T E(\mu))\mu$. (Note that $G(\mu)\mu$ is a matrix-vector product.)*

Proof of Claim. This is a straightforward calculation. Consider the i -th coordinate of $\nabla F(\mu)$, i.e., the i -th partial derivative

$$\begin{aligned} \frac{\partial F}{\partial \mu_i} &= \frac{\partial(\sum_j \mu_j E(\mu)_j)}{\partial \mu_i} \\ &= E(\mu)_i + \sum_{j=1}^d \mu_j \frac{\partial E(\mu)_j}{\partial \mu_i} \\ &= E(\mu)_i + \sum_{j=1}^d \mu_j \left(\int_{\mathbb{R}^d} (\frac{n}{2\pi})^{n/2} (\eta(y) - \eta^*(y))_j \frac{\partial(e^{-\frac{n\|y-\mu\|^2}{2}})}{\partial \mu_i} dy \right) \\ &= E(\mu)_i + \mu^T \int_{\mathbb{R}^d} (\frac{n}{2\pi})^{n/2} (\eta(y) - \eta^*(y)) \frac{\partial(e^{-\frac{n\|y-\mu\|^2}{2}})}{\partial \mu_i} dy \\ &= E(\mu)_i + \mu^T \int_{\mathbb{R}^d} (\frac{n}{2\pi})^{n/2} (\eta(y) - \eta^*(y)) e^{-\frac{n\|y-\mu\|^2}{2}} (n(y_i - \mu_i)) dy \\ &= E(\mu)_i + n\mu^T G^i(\mu) - n(\mu^T E(\mu))\mu_i \end{aligned}$$

where in the third equality, we have used Theorem 2.4 and the fact that $\eta(y), \eta^*(y)$ do not depend on μ by sufficiency of the sample average (Proposition 2.3). The last expression above corresponds to the i -th coordinate of $E(\mu) + nG(\mu)\mu - n(\mu^T E(\mu))\mu$. Thus, we are done. \diamond

CLAIM 3.3. $\nabla^2 F(0) = n(G(0)^T + G(0))$.

Proof of Claim. Let us compute $\frac{\partial^2 F}{\partial \mu_i \partial \mu_j}$ using the expression for $\frac{\partial F}{\partial \mu_i}$ from Claim 3.2.

$$\begin{aligned} \frac{\partial^2 F}{\partial \mu_i \partial \mu_j} &= \frac{\partial(E(\mu)_i)}{\partial \mu_j} + n \frac{\partial(\mu^T G^i(\mu))}{\partial \mu_j} - n \frac{\partial((\mu^T E(\mu))\mu_i)}{\partial \mu_j} \\ &= \frac{\partial(E(\mu)_i)}{\partial \mu_j} + n(G^i(\mu))_j + n\mu^T \frac{\partial(G^i)}{\partial \mu_j} - n \frac{\partial(\mu^T E(\mu))}{\partial \mu_j} \mu_i - n(\mu^T E(\mu))\gamma_{ij} \end{aligned}$$

where γ_{ij} denotes the Kronecker delta function, i.e., $\gamma_{ij} = 1$ if $i = j$ and 0 otherwise. At $\mu = 0$, the above simplifies to

$$\frac{\partial^2 F}{\partial \mu_i \partial \mu_j} \Big|_{\mu=0} = \frac{\partial(E(\mu)_i)}{\partial \mu_j} \Big|_{\mu=0} + n(G^i(0))_j. \quad (3.2)$$

Let us now investigate $\frac{\partial(E(\mu)_i)}{\partial\mu_j}$. By applying Theorem 2.4 and the sufficiency of $\bar{\xi}$ again, we obtain

$$\begin{aligned}\frac{\partial(E(\mu)_i)}{\partial\mu_j} &= \int_{\mathbb{R}^d} \left(\frac{n}{2\pi}\right)^{n/2} (\eta(y) - \eta^*(y))_i \frac{\partial(e^{\frac{-n\|y-\mu\|^2}{2}})}{\partial\mu_j} dy \\ &= \int_{\mathbb{R}^d} \left(\frac{n}{2\pi}\right)^{n/2} (\eta(y) - \eta^*(y))_i e^{\frac{-n\|y-\mu\|^2}{2}} (n(y_j - \mu_j)) dy \\ &= n\left(\frac{n}{2\pi}\right)^{n/2} \int_{\mathbb{R}^d} y_j (\eta(y) - \eta^*(y))_i e^{\frac{-n\|y-\mu\|^2}{2}} dy \\ &\quad - n\left(\frac{n}{2\pi}\right)^{n/2} \mu_j \int_{\mathbb{R}^d} (\eta(y) - \eta^*(y))_i e^{\frac{-n\|y-\mu\|^2}{2}} dy \\ &= n(G^j(\mu))_i - n\mu_j E(\mu)_i\end{aligned}$$

Therefore, at $\mu = 0$, we obtain that $\frac{\partial(E(\mu)_i)}{\partial\mu_j} \Big|_{\mu=0} = nG^j(0)_i$. Putting this back into (3.2), and using the definition of the matrix $G(\mu)$, we obtain

$$\frac{\partial^2 F}{\partial\mu_i\mu_j} \Big|_{\mu=0} = nG^j(0)_i + n(G^i(0))_j = n(G(0)_{ji} + G(0)_{ij}).$$

Thus, we obtain that $\nabla^2 F(0) = n(G(0)^T + G(0))$. \diamond

CLAIM 3.4. *There exists a direction of positive curvature for $\nabla^2 F(0)$, i.e., there exists $d \in \mathbb{R}^d$ such that $d^T \nabla^2 F(0)d > 0$.*

Proof of Claim. Consider the trace $\text{Tr}(\nabla^2 F(0))$ of the Hessian at $\mu = 0$. By Claim 3.3,

$$\begin{aligned}\text{Tr}(\nabla^2 F(0)) &= 2n \text{Tr}(G(0)) \\ &= 2n \sum_{i=1}^d \left(\frac{n}{2\pi}\right)^{n/2} \int_{\mathbb{R}^d} y_i (\eta(y) - \eta^*(y))_i e^{\frac{-n\|y-\mu\|^2}{2}} dy \\ &= 2n\left(\frac{n}{2\pi}\right)^{n/2} \int_{\mathbb{R}^d} y^T (\eta(y) - \eta^*(y)) e^{\frac{-n\|y-\mu\|^2}{2}} dy\end{aligned}$$

By Claim 3.1, $y^T (\eta(y) - \eta^*(y)) \geq 0$ for any $y \in \mathbb{R}^d$ and $y^T (\eta(y) - \eta^*(y)) > 0$ on a set of strictly positive measure. Therefore, $\int_{\mathbb{R}^d} y^T (\eta(y) - \eta^*(y)) e^{\frac{-n\|y\|^2}{2}} dy > 0$.

Therefore, the trace of $\nabla^2 F(0)$ is strictly positive. Since the trace equals the sum of the eigenvalues of $\nabla^2 F(0)$ (see Section 1.2.5 in [19]), we must have at least one strictly positive eigenvalue. The corresponding eigenvector is a direction of positive curvature. \diamond

As noted earlier, $F(0) = 0$. Combining Claim 3.4 and Lemma 2.1, there exists $\bar{\mu} \in \mathbb{R}^d$ such that $F(\bar{\mu}) > 0$. \square

3.2 General covariance

The proof in the previous section focused on the family of normal distributions with the identity as the covariance matrix. We now consider any positive definite covariance matrix Σ for the normal distribution of ξ . In this case, we again consider the function $F(\mu)$ defined in (3.1) and prove that there exists $\bar{\mu}$ such that $F(\bar{\mu}) > 0$. The only difference is that in

the formulas one must substitute the distribution $\bar{\xi} \sim N(\mu, \frac{1}{n}\Sigma)$, i.e., the density function everywhere must be

$$g_{\mu, \Sigma}(y) := \frac{1}{\sqrt{\sigma}} \left(\frac{n}{2\pi} \right)^{n/2} \exp \left(-\frac{n}{2}(y - \mu)^T \Sigma^{-1}(y - \mu) \right),$$

where σ is the determinant of Σ . Redefining

$$\begin{aligned} E(\mu) &:= \frac{1}{\sqrt{\sigma}} \left(\frac{n}{2\pi} \right)^{n/2} \int_{\mathbb{R}^d} (\eta(y) - \eta^*(y)) \exp \left(-\frac{n}{2}(y - \mu)^T \Sigma^{-1}(y - \mu) \right) dy, \\ G^i(\mu) &:= \frac{1}{\sqrt{\sigma}} \left(\frac{n}{2\pi} \right)^{n/2} \int_{\mathbb{R}^d} y_i (\eta(y) - \eta^*(y)) \exp \left(-\frac{n}{2}(y - \mu)^T \Sigma^{-1}(y - \mu) \right) dy, \quad i = 1, \dots, d, \end{aligned}$$

letting $G(\mu)$ be the matrix with $G^i(\mu)$ as rows, and adapting the calculations from the previous section reveals that

$$\nabla^2 F(0) = n(\Sigma^{-1}G(0) + G(0)^T \Sigma^{-1}) \quad (3.3)$$

Claim 3.1 again shows that the trace $\text{Tr}(G(0)) = \int_{\mathbb{R}^d} y^T (\eta(y) - \eta^*(y)) g_{0, \Sigma}(y) dy > 0$. This shows that $G(0)$ has an eigenvalue λ with positive *real* part (since $G(0)$ is not guaranteed to be symmetric, its eigenvalues and eigenvectors may be complex). Let the corresponding (possibly complex) eigenvector be v , i.e., $G(0)v = \lambda v$ and $\text{Re}(\lambda) > 0$ (denoting the real part of λ). Following standard linear algebra notation, for any matrix/vector M , M^* will denote its Hermitian conjugate [19] (which equals the transpose if the matrix has real entries). We now consider

$$\begin{aligned} v^* \nabla^2 F(0) v &= n(v^*(\Sigma^{-1}G(0) + G(0)^T \Sigma^{-1})v) \\ &= n(v^* \Sigma^{-1}G(0)v + v^* G(0)^T \Sigma^{-1}v) \\ &= n(v^* \Sigma^{-1}G(0)v + v^* G(0)^* \Sigma^{-1}v) \\ &= n(\lambda(v^* \Sigma^{-1}v) + \lambda^*(v^* \Sigma^{-1}v)) \\ &= 2n(v^* \Sigma^{-1}v) \text{Re}(\lambda) \end{aligned}$$

Since Σ is positive definite, so is Σ^{-1} . Therefore $v^* \Sigma^{-1}v > 0$ and we obtain that $v^* \nabla^2 F(0)v > 0$. Since $\nabla^2 F(0)$ is a symmetric matrix, all its eigenvalues are real and in particular its largest eigenvalue γ_d is positive because

$$\gamma_d = \max_{x \in \mathbb{C}^d \setminus \{0\}} \frac{x^* \nabla^2 F(0)x}{x^* x} \geq \frac{v^* \nabla^2 F(0)v}{v^* v} > 0.$$

Thus, $\nabla^2 F(0)$ has a direction of positive curvature and Lemma 2.1 implies that there exists $\bar{\mu} \in \mathbb{R}^d$ such that $F(\bar{\mu}) > 0$.

4 An alternate proof for the linear objective based on Bayes' decision rules

To the best of our knowledge, our proof technique for admissibility from the previous sections is new. The conventional way of addressing admissibility uses Bayesian analysis. We recall

the basic ideas behind these techniques and provide an alternate proof for Theorem 1.1 using these ideas, which arguably gives a simpler proof. On the other hand, this alternate proof builds upon some well-established facts in statistics, and so is less of a “first principles” proof compared to the one presented in the previous sections. Moreover, as we noted earlier, the new technique of the previous proof might be useful for future admissibility investigations in stochastic optimization.

We now briefly review the relevant ideas from Bayesian analysis. Consider a general statistical decision problem with Θ denoting the states of nature, \mathcal{A} denoting the set of actions, and $\{P_\theta : \theta \in \Theta\}$ the family of distributions on the sample space χ . Let P^* be any so-called *prior distribution* on Θ . For any decision rule δ , one can compute the expected risk, a.k.a., the *Bayes’ risk*

$$r(P^*, \delta) := \mathbb{E}_{\theta \sim P^*}[R(\theta, \delta)].$$

A decision rule that minimizes $r(P^*, \delta)$ is said to be a *Bayes’ decision rule*.

THEOREM 4.1. [24, Chapter 5, Theorem 2.4] *If a decision rule is the unique³ Bayes’ decision rule for some prior, then it is admissible.*

The following is the well-known statement that Gaussian distributions are *self-conjugate* [24, Example 2.2].

THEOREM 4.2. *Let $d \in \mathbb{N}$ and let $\Sigma \in \mathbb{R}^{d \times d}$ be fixed. For the joint distribution on $(\xi, \mu) \in \underbrace{(\mathbb{R}^d \times \mathbb{R}^d \times \dots \times \mathbb{R}^d)}_{ntimes} \times \mathbb{R}^d$ defined by $\xi|\mu \sim \underbrace{\mathcal{N}(\mu, \Sigma) \times \mathcal{N}(\mu, \Sigma) \times \dots \times \mathcal{N}(\mu, \Sigma)}_{ntimes}$ and $\mu \sim \mathcal{N}(\mu_0, \Sigma_0)$, we have that $\mu|\xi \sim \mathcal{N}((\Sigma_0^{-1} + n\Sigma^{-1})^{-1}(\Sigma_0^{-1}\mu_0 + n\Sigma^{-1}\bar{\xi}), (\Sigma_0^{-1} + n\Sigma^{-1})^{-1})$.*

Alternate proof of Theorem 1.1. Consider the prior P^* to be $\mu \sim \mathcal{N}(0, \Sigma)$, then by Theorem 4.2, $\mu|\xi \sim \mathcal{N}(\frac{n}{n+1}\bar{\xi}, \frac{1}{n+1}\Sigma)$. In particular, the mean of μ , conditioned on the observation ξ is simply a scaling of the sample average $\bar{\xi}$. Now we do a standard Bayesian analysis:

$$\begin{aligned} r(P^*, \delta) &= \mathbb{E}_{\mu \sim \mathcal{N}(0, \Sigma)}[\mathbb{E}_{\xi \sim \mathcal{N}(\mu, \Sigma)^n}[\mathcal{L}(\mu, \delta)]] \\ &= \int_{\mu} \int_{\xi} (\mu^T \delta(\xi) - \mu^T x(\mu)) p(\xi|\mu) p(\mu) d\xi d\mu \\ &= \int_{\mu} \int_{\xi} \mu^T \delta(\xi) p(\xi|\mu) p(\mu) d\xi d\mu - C \end{aligned}$$

where $x(\mu)$ again denotes the optimal solution to $\min\{\mu^T x : x \in X\}$, $p(\xi|\mu)$ denotes the conditional density function of $\xi|\mu$, $p(\mu)$ is the density function of the prior on μ , and the constant C equals $\int_{\mu} \int_{\xi} \mu^T x(\mu) p(\xi|\mu) p(\mu) d\xi d\mu = \int \mu^T x(\mu) p(\mu) d\mu$. To find the decision rule δ that minimizes $r(P^*, \delta)$, we thus need to minimize $\int_{\mu} \int_{\xi} \mu^T \delta(\xi) p(\xi|\mu) p(\mu) d\xi d\mu$. We change the order of integration by Fubini’s theorem, and rewrite

$$r(P^*, \delta) = \int_{\xi} \int_{\mu} \mu^T \delta(\xi) p(\mu|\xi) p(\xi) d\mu d\xi.$$

³Here uniqueness is to be interpreted up to differences on a set of measure zero.

Consequently, given the observation ξ , we choose $\delta(\xi) \in X$ that minimizes the inner integral

$$\int_{\mu} \mu^T \delta(\xi) p(\mu|\xi) d\mu = \frac{n}{n+1} \bar{\xi}^T \delta(\xi).$$

Thus, we may set $\delta(\xi)$ to be the minimizer in X for the linear objective vector $\frac{n}{n+1} \bar{\xi}^T$, which is just a scaling of the sample average. Except for a set of measure zero, any linear objective has a unique solution as was noted in the proof of Claim 3.1. Thus, the Bayes' decision rule is unique and coincides with the sample average estimator δ_{SA}^n . We are done by appealing to Theorem 4.1. \square

5 Proof of Theorem 1.2 (the scenario with quadratic objective)

A more general version of Theorem 4.1 goes by the name of *Blyth's method*. Here, we state it as in [24] (see Exercise 7.12 in Chapter 5).

THEOREM 5.1. *Let $\Theta \subseteq \mathbb{R}^d$ be any open set of states of nature. Suppose δ is a decision rule with a continuous risk function and $\{\pi_n\}_{n \in \mathbb{N}}$ is a sequence of prior distributions such that:*

1. $r(\pi_n, \delta) = \int \mathcal{R}(\theta, \delta) d\pi_n < \infty$ for all $n \in \mathbb{N}$, where r is the Bayes risk.
2. For any nonempty open subset $\Theta_0 \subseteq \Theta$, we have

$$\lim_{n \rightarrow \infty} \frac{r(\pi_n, \delta) - r(\pi_n, \delta^{\pi_n})}{\int_{\Theta_0} \pi_n(\theta) d\theta} = 0$$

where δ^{π_n} is a Bayes decision rule having finite Bayes risk with respect to the prior density π_n . Then, δ is an admissible decision rule.

Proof of Theorem 1.2. For simplicity of exposition, we consider X to be centered at 0, i.e., $\ell_i = -u_i$. The entire proof can be reproduced for the general case by translating the means of the priors to the center of axis-aligned box X . The calculations are much easier to read and follow when we assume the origin to be the center. We will show that δ_{SA}^n satisfies the conditions of Theorem 5.1 under the priors

$$\mu \sim \mathcal{N}(0, \tau^2 \Sigma), \quad \tau \in \mathbb{N}, \text{ and } \Sigma \text{ is the covariance matrix of } \xi | \mu.$$

First we obtain a simple expression for the loss function $\mathcal{L}(\mu, x)$ for any action $x \in X$ under the state of nature $\mu \in \mathbb{R}^d$. As noted in Section 1.2, $\mathbb{E}_{\xi} [F(x, \xi)] = \mathbb{E}_{\xi} [\frac{1}{2} x^T x - \xi^T x] = \frac{1}{2} x^T x - \mu^T x = \frac{1}{2} \|\mu - x\|^2 - \frac{1}{2} \|\mu\|^2$. Let the minimum value of this for $x \in X$ be denoted by $B(\mu)$. Thus,

$$\begin{aligned} \mathcal{L}(\mu, x) &= \mathbb{E}_{\xi \sim \mathcal{N}(\mu, \Sigma)} [F(x, \xi)] - \mathbb{E}_{\xi \sim D} [F(x(D), \xi)] \\ &= \frac{1}{2} \|\mu - x\|^2 - \frac{1}{2} \|\mu\|^2 - B(\mu) \end{aligned}$$

Putting this into the numerator of the second condition in Theorem 5.1 and simplifying (which means that the terms $-\frac{1}{2} \|\mu\|^2 - B(\mu)$ cancel out), we need to show that for any open set $\Omega_0 \subseteq \mathbb{R}^d$,

$$\lim_{\tau \rightarrow \infty} \frac{\int_{\mathbb{R}^d} \int_{\mathbb{R}^d} \|\mu - \delta_{SA}^n(\xi)\|^2 p(\xi|\mu) \pi_\tau(\mu) d\xi d\mu - \int_{\mathbb{R}^d} \int_{\mathbb{R}^d} \|\mu - \delta_{Bayes}(\xi)\|^2 p(\xi|\mu) \pi_\tau(\mu) d\xi d\mu}{\int_{\Omega_0} \pi_\tau(\mu) d\mu} = 0 \quad (5.1)$$

where $p(\xi|\mu) = \mathcal{N}(\mu, \Sigma)$ denotes the conditional density of ξ given μ , and $\pi_\tau(\mu)$ is the marginal density of μ (of course, the marginal density $\pi_\tau(\mu)$ is nothing but the prior $\mathcal{N}(0, \tau^2 \Sigma)$).

Next, let us see how the rule δ_{SA}^n given by (1.5) behaves. Minimizing $\frac{1}{2}\|x\|^2 - \bar{\xi}^T x$ is equivalent to minimizing $\frac{1}{2}\|x - \bar{\xi}\|^2$ since $\bar{\xi}$ can be regarded as constant for the optimization problem $\min_{x \in X} \frac{1}{2}\|x\|^2 - \bar{\xi}^T x$. Thus, δ_{SA}^n returns the closest point to $\bar{\xi}$ in X , i.e.,

$$\delta_{SA}^n(\xi) = \text{Proj}_X(\bar{\xi}) \quad (5.2)$$

where the notation $\text{Proj}_X(y)$ denotes the projection of the closest point in X to y .

Let us also see what the Bayes' rule δ_{Bayes} is, i.e., what value of $\delta_{Bayes}(\xi)$ minimizes

$$\begin{aligned} \int_{\mathbb{R}^d} \int_{\mathbb{R}^d} \|\mu - \delta_{SA}^n(\xi)\|^2 p(\xi|\mu) \pi_\tau(\mu) d\xi d\mu &= \int_{\mathbb{R}^d} \int_{\mathbb{R}^d} \|\mu - \delta_{Bayes}(\xi)\|^2 p(\mu|\xi) m(\xi) d\mu d\xi \\ &= \int_{\mathbb{R}^d} \mathbb{E}_{\mu|\xi} \|\mu - \delta_{Bayes}(\xi)\|^2 m(\xi) d\xi. \end{aligned}$$

where we have again evaluated the Bayes' risks by switching the order of the integrals and using the conditional density

$$p(\mu|\xi) = \mathcal{N}\left(\frac{n\tau^2}{n\tau^2 + 1} \bar{\xi}, \frac{\tau^2}{n\tau^2 + 1} \Sigma\right) \quad (\text{by Theorem 4.2}), \quad (5.3)$$

and the marginal density of ξ is denoted by $m(\xi)$. Since

$$\mathbb{E}[\|Y - a\|^2] = \|\mathbb{E}[Y] - a\|^2 + \mathbb{V}[Y] \quad (5.4)$$

for any random variable Y and constant $a \in \mathbb{R}$, one sees that

$$\delta_{Bayes}(\xi) = \text{Proj}_X\left(\frac{n\tau^2}{n\tau^2 + 1} \bar{\xi}\right). \quad (5.5)$$

Let us now consider the numerator and denominator of the left hand side in (5.1) separately.

Numerator: The numerator is:

$$\begin{aligned} &\int_{\mathbb{R}^d} \int_{\mathbb{R}^d} \|\mu - \delta_{SA}^n(\xi)\|^2 p(\xi|\mu) \pi_\tau(\mu) d\xi d\mu - \int_{\mathbb{R}^d} \int_{\mathbb{R}^d} \|\mu - \delta_{Bayes}(\xi)\|^2 p(\xi|\mu) \pi_\tau(\mu) d\xi d\mu \\ &= \int_{\mathbb{R}^d} \int_{\mathbb{R}^d} \|\mu - \text{Proj}_X(\bar{\xi})\|^2 p(\xi|\mu) \pi_\tau(\mu) d\xi d\mu - \int_{\mathbb{R}^d} \int_{\mathbb{R}^d} \left\|\mu - \text{Proj}_X\left(\frac{n\tau^2}{n\tau^2 + 1} \bar{\xi}\right)\right\|^2 p(\xi|\mu) \pi_\tau(\mu) d\xi d\mu \\ &= \int_{\mathbb{R}^d} \int_{\mathbb{R}^d} \int_{\mathbb{R}^d} \|\mu - \text{Proj}_X(\bar{\xi})\|^2 p(\xi|\bar{\xi}) p(\bar{\xi}|\mu) \pi_\tau(\mu) d\xi d\bar{\xi} d\mu \\ &\quad - \int_{\mathbb{R}^d} \int_{\mathbb{R}^d} \int_{\mathbb{R}^d} \left\|\mu - \text{Proj}_X\left(\frac{n\tau^2}{n\tau^2 + 1} \bar{\xi}\right)\right\|^2 p(\xi|\bar{\xi}) p(\bar{\xi}|\mu) \pi_\tau(\mu) d\xi d\bar{\xi} d\mu \\ &= \int_{\mathbb{R}^d} \int_{\mathbb{R}^d} \|\mu - \text{Proj}_X(\bar{\xi})\|^2 p(\bar{\xi}|\mu) \pi_\tau(\mu) d\bar{\xi} d\mu - \int_{\mathbb{R}^d} \int_{\mathbb{R}^d} \left\|\mu - \text{Proj}_X\left(\frac{n\tau^2}{n\tau^2 + 1} \bar{\xi}\right)\right\|^2 p(\bar{\xi}|\mu) \pi_\tau(\mu) d\bar{\xi} d\mu \\ &= \int_{\mathbb{R}^d} \int_{\mathbb{R}^d} \|\mu - \text{Proj}_X(\bar{\xi})\|^2 p(\mu|\bar{\xi}) m(\bar{\xi}) d\mu d\bar{\xi} - \int_{\mathbb{R}^d} \int_{\mathbb{R}^d} \left\|\mu - \text{Proj}_X\left(\frac{n\tau^2}{n\tau^2 + 1} \bar{\xi}\right)\right\|^2 p(\mu|\bar{\xi}) m(\bar{\xi}) d\mu d\bar{\xi}, \end{aligned}$$

where the first equality follows from substituting (5.2) and (5.5), and the last equality follows from the standard trick in Bayesian analysis of switching the order of the integrals.

Since $\bar{\xi}|\mu \sim \mathcal{N}(\mu, \frac{1}{n}\Sigma)$, and $\pi(\mu) \sim \mathcal{N}(0, \tau^2\Sigma)$, it is a simple exercise to check that $\mu|\bar{\xi} \sim \mathcal{N}(\frac{n\tau^2}{n\tau^2+1}\bar{\xi}, \frac{\tau^2}{n\tau^2+1}\Sigma)$, and $m(\bar{\xi}) \sim \mathcal{N}(0, \frac{n\tau^2+1}{n}\Sigma)$. The formula for the numerator above can then be rewritten as

$$\begin{aligned}
& \int_{\mathbb{R}^d} \mathbb{E}_{\mu|\bar{\xi}} \left\| \mu - \text{Proj}_X(\bar{\xi}) \right\|^2 m(\bar{\xi}) d\bar{\xi} - \int_{\mathbb{R}^d} \mathbb{E}_{\mu|\bar{\xi}} \left\| \mu - \text{Proj}_X\left(\frac{n\tau^2}{n\tau^2+1}\bar{\xi}\right) \right\|^2 m(\bar{\xi}) d\bar{\xi} \\
&= \int_{\mathbb{R}^d} \mathbb{V}(\mu|\bar{\xi}) + \left\| \mathbb{E}_{\mu|\bar{\xi}} \mu - \text{Proj}_X(\bar{\xi}) \right\|^2 m(\bar{\xi}) d\bar{\xi} - \int_{\mathbb{R}^d} \mathbb{V}(\mu|\bar{\xi}) + \left\| \mathbb{E}_{\mu|\bar{\xi}} \mu - \text{Proj}_X\left(\frac{n\tau^2}{n\tau^2+1}\bar{\xi}\right) \right\|^2 m(\bar{\xi}) d\bar{\xi} \\
&= \int_{\mathbb{R}^d} \left\| \frac{n\tau^2}{n\tau^2+1}\bar{\xi} - \text{Proj}_X(\bar{\xi}) \right\|^2 - \left\| \frac{n\tau^2}{n\tau^2+1}\bar{\xi} - \text{Proj}_X\left(\frac{n\tau^2}{n\tau^2+1}\bar{\xi}\right) \right\|^2 m(\bar{\xi}) d\bar{\xi} \\
&= \int_{\mathbb{R}^d} \left(\left\| \text{Proj}_X\left(\frac{n\tau^2}{n\tau^2+1}\bar{\xi}\right) - \text{Proj}_X(\bar{\xi}) \right\|^2 - 2 \left\langle \frac{n\tau^2}{n\tau^2+1}\bar{\xi} - \text{Proj}_X\left(\frac{n\tau^2}{n\tau^2+1}\bar{\xi}\right), \text{Proj}_X(\bar{\xi}) - \text{Proj}_X\left(\frac{n\tau^2}{n\tau^2+1}\bar{\xi}\right) \right\rangle \right) m(\bar{\xi}) d\bar{\xi}, \tag{5.6}
\end{aligned}$$

where the first equality follows from (5.4), the second equality follows from the formula for the conditional density $p(\mu|\bar{\xi})$, and the last equality follows from the fact that

$$\|a - b\|^2 - \|a - c\|^2 = \|c - b\|^2 - 2\langle a - c, b - c \rangle, \quad \text{for any three vectors } a, b, c \in \mathbb{R}^d.$$

CLAIM 5.2. Consider the box $X = \{x \in \mathbb{R}^d : -u_i \leq x_i \leq u_i \quad i = 1, \dots, d\}$. Let F be any face of X and so $F' := \frac{n\tau^2+1}{n\tau^2}F$ is a face of $X' := \frac{n\tau^2+1}{n\tau^2}X$. Suppose $\bar{\xi} \in F' + N_{F'}$ where $N_{F'}$ denotes the normal cone at F' with respect to X' (see Lemma 2.5 and the discussion above it), then $\text{Proj}_X\left(\frac{n\tau^2}{n\tau^2+1}\bar{\xi}\right)$ and $\text{Proj}_X(\bar{\xi})$ both lie in F . Consequently,

$$\left\langle \frac{n\tau^2}{n\tau^2+1}\bar{\xi} - \text{Proj}_X\left(\frac{n\tau^2}{n\tau^2+1}\bar{\xi}\right), \text{Proj}_X(\bar{\xi}) - \text{Proj}_X\left(\frac{n\tau^2}{n\tau^2+1}\bar{\xi}\right) \right\rangle = 0. \tag{5.7}$$

Proof. The general formula for the projection onto X is given by

$$\text{Proj}_X(y) = p, \quad \text{where } p_i = \text{sign}(y_i) \min\{|y_i|, u_i\}. \tag{5.8}$$

Consider any face F of X and any $\bar{\xi} \in F' + N_{F'}$; see Figure 1. By Lemma 2.5 part 1., $|\bar{\xi}_i| \geq \frac{n\tau^2+1}{n\tau^2}|u_i|$ for all $i \in I_F^+ \cup I_F^-$. Therefore, $|\bar{\xi}_i| \geq \frac{n\tau^2}{n\tau^2+1}|\bar{\xi}_i| \geq |u_i|$ for all $i \in I_F^+ \cup I_F^-$. By the projection formula (5.8), for all $i \in I_F^+ \cup I_F^-$ the i -th coordinate of both $\text{Proj}_X\left(\frac{n\tau^2}{n\tau^2+1}\bar{\xi}\right)$ and $\text{Proj}_X(\bar{\xi})$ are equal to $\text{sign}(\bar{\xi}_i)u_i$. This shows that they both lie on F . By the geometry of projections, the vector $\frac{n\tau^2}{n\tau^2+1}\bar{\xi} - \text{Proj}_X\left(\frac{n\tau^2}{n\tau^2+1}\bar{\xi}\right)$ is orthogonal to the face F that contains the projection $\text{Proj}_X\left(\frac{n\tau^2}{n\tau^2+1}\bar{\xi}\right)$. This proves (5.7). \square

By appealing to (5.7), one can reduce (5.6) to

$$\int_{\mathbb{R}^d} \left\| \text{Proj}_X\left(\frac{n\tau^2}{n\tau^2+1}\bar{\xi}\right) - \text{Proj}_X(\bar{\xi}) \right\|^2 m(\bar{\xi}) d\bar{\xi} \tag{5.9}$$

Using Lemma 2.5 part 2., we decompose the above integral as follows:

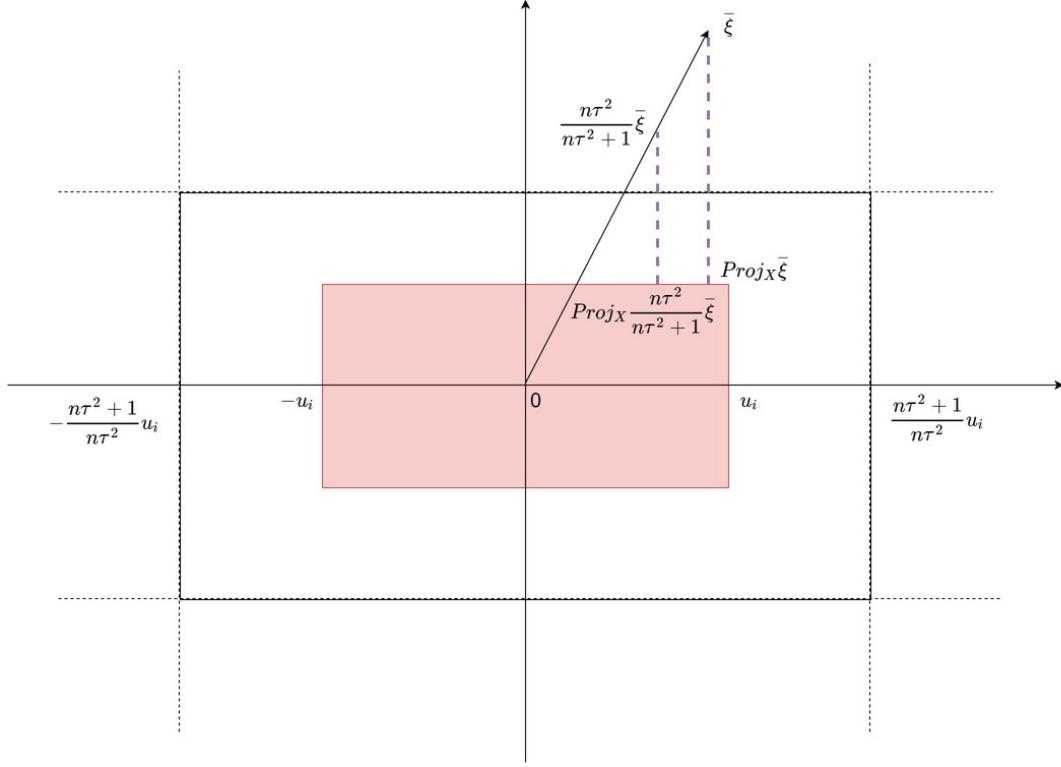


Figure 1: The scaling X' of X and the different regions $F' + N_{F'}$ for different faces F of X , with an illustration of Claim 5.2.

$$(5.9) = \sum_{F \text{ face of } X} \int_{F' + N_{F'}} \left\| \text{Proj}_X \left(\frac{n\tau^2}{n\tau^2+1} \bar{\xi} \right) - \text{Proj}_X(\bar{\xi}) \right\|^2 m(\bar{\xi}) d\bar{\xi}.$$

where we have used the notation of Claim 5.2 for F' . Using the formula in Lemma 2.5 part 1., we may simplify the integral further by introducing some additional notation. For any face F of X , recall the notation I_F^+ and I_F^- from Lemma 2.5. Let $I_F^0 := \{1, \dots, d\} \setminus (I_F^+ \cup I_F^-)$. We introduce the decomposition of any vector $y \in \mathbb{R}^d$ into y^+, y^-, y^0 where $y^+ \in \mathbb{R}^{I_F^+}$ denotes the restriction of the vector y onto the coordinates in I_F^+ ; similarly, $y^- \in \mathbb{R}^{I_F^-}$ is y restricted to I_F^- , and y^0 is y restricted to I_F^0 . Denote the corresponding domains $D_F^+ := \{z \in \mathbb{R}^{I_F^+} : z_i \geq \frac{n\tau^2+1}{n\tau^2}u_i\}$, $D_F^- := \{z \in \mathbb{R}^{I_F^-} : z_i \leq -\frac{n\tau^2+1}{n\tau^2}u_i\}$ and $D_F^0 := \{z \in \mathbb{R}^{I_F^0} : -\frac{n\tau^2+1}{n\tau^2}u_i \leq z_i \leq \frac{n\tau^2+1}{n\tau^2}u_i\}$. By Lemma 2.5 part 1.,

$$\begin{aligned}
(5.9) &= \sum_{F \text{ face of } X} \int_{D_F^0} \int_{D_F^-} \int_{D_F^+} \left\| \text{Proj}_X \left(\frac{n\tau^2}{n\tau^2+1} \bar{\xi} \right) - \text{Proj}_X(\bar{\xi}) \right\|^2 m(\bar{\xi}) d\bar{\xi}^+ d\bar{\xi}^- d\bar{\xi}^0 \\
&\leq \sum_{F \text{ face of } X} \int_{D_F^0} \int_{D_F^-} \int_{D_F^+} \left(\frac{1}{(n\tau^2+1)^2} \sum_{i \in I_F^0} \bar{\xi}_i^2 \right) m(\bar{\xi}) d\bar{\xi}^+ d\bar{\xi}^- d\bar{\xi}^0
\end{aligned}$$

where the inequality follows from Claim 5.2 which tells us that both $\text{Proj}_X\left(\frac{n\tau^2}{n\tau^2+1}\bar{\xi}\right)$ and $\text{Proj}_X(\bar{\xi})$ lie on F and therefore coincide on the coordinates in $I_F^+ \cup I_F^-$; thus, the coordinates in $I_F^+ \cup I_F^-$ vanish in the integrand. Moreover, on any remaining coordinate i that is not set to the bound u_i or $-u_i$, the absolute difference in the coordinate is at most $|\frac{n\tau^2}{n\tau^2+1}\bar{\xi}_i - \bar{\xi}_i| = \frac{1}{n\tau^2+1}|\bar{\xi}_i|$, by the projection formula (5.8). Plugging in the formula for $m(\bar{\xi})$, we get

$$\begin{aligned}
(5.9) &\leq \left(\frac{n}{2\pi(n\tau^2+1)\det(\Sigma)^{1/d}}\right)^{d/2} \sum_{F \text{ face of } X} \int_{D_F^0} \int_{D_F^-} \int_{D_F^+} \left(\frac{1}{(n\tau^2+1)^2} \sum_{i \in I_F^0} \bar{\xi}_i^2\right) e^{\frac{-n\bar{\xi}^T \Sigma^{-1} \bar{\xi}}{2(n\tau^2+1)}} d\bar{\xi}^+ d\bar{\xi}^- d\bar{\xi}^0 \\
&= \left(\frac{n}{2\pi(n\tau^2+1)\det(\Sigma)^{1/d}}\right)^{d/2} \sum_{F \text{ face of } X} \int_{D_F^0} \left(\frac{1}{(n\tau^2+1)^2} \sum_{i \in I_F^0} \bar{\xi}_i^2\right) \int_{D_F^+} \int_{D_F^-} e^{\frac{-n\bar{\xi}^T \Sigma^{-1} \bar{\xi}}{2(n\tau^2+1)}} d\bar{\xi}^+ d\bar{\xi}^- d\bar{\xi}^0 \\
&= \left(\frac{n}{2\pi(n\tau^2+1)\det(\Sigma)^{1/d}}\right)^{d/2} \sum_{F \text{ face of } X} \int_{D_F^0} \left(\frac{1}{(n\tau^2+1)^2} \sum_{i \in I_F^0} \bar{\xi}_i^2\right) \cdot C_F (n\tau^2+1)^{(d-\dim(F))/2} h(\bar{\xi}^0, \tau) d\bar{\xi}^0
\end{aligned}$$

where the integral is evaluated using standard Gaussian integral formulas and C_F is a constant independent of τ , $\dim(F)$ denotes the dimension of the face F and $h(\bar{\xi}^0, \tau)$ is a continuous function which is upper bounded on the compact domain D_F^0 for every F by a universal constant independent of τ .

We now observe that if $I_F^0 = \emptyset$, i.e., F is a vertex of X , then the integrand is simply 0 (in other words, when $\bar{\xi}$ lies in the normal cone of a vertex v of X translated by $\frac{n\tau^2+1}{n\tau^2}v$, the projections $\text{Proj}_X\left(\frac{n\tau^2}{n\tau^2+1}\bar{\xi}\right)$ and $\text{Proj}_X(\bar{\xi})$ are both equal to v , and the integral vanishes). Therefore, we are left with the terms where the face has dimensional least 1. Thus, $(n\tau^2+1)^{(d-\dim(F))/2}$ can be upper bounded by $(n\tau^2+1)^{(d-1)/2}$. Since D_F^0 is a compact domain and h is continuous function upper bounded by a universal constant independent of τ , we infer the following upper bound on the numerator

$$\left(\frac{n}{2\pi(n\tau^2+1)\det(\Sigma)^{1/d}}\right)^{d/2} \cdot C \cdot (n\tau^2+1)^{(d-5)/2}$$

where C is a constant independent of τ .

Denominator: Using the density formula for $\pi_\tau(\mu)$, the denominator of (5.1) is

$$\frac{1}{[2\pi\tau^2]^{d/2} \det(\Sigma)^{1/2}} \int_{\Omega_0} e^{\frac{-\mu^T \Sigma^{-1} \mu}{2\tau^2}} d\mu.$$

Combining the formulas for the numerator and denominator in (5.1), we have:

$$\begin{aligned}
&\lim_{\tau \rightarrow \infty} \frac{\int_{\mathbb{R}^d} \int_{\mathbb{R}^d} \|\mu - \delta_{SA}^n(\xi)\|^2 p(\xi|\mu) \pi_\tau(\mu) d\xi d\mu - \int_{\mathbb{R}^d} \int_{\mathbb{R}^d} \|\mu - \delta_{Bayes}(\xi)\|^2 p(\xi|\mu) \pi_\tau(\mu) d\xi d\mu}{\int_{\Omega_0} \pi_\tau(\xi) d\xi} \\
&\leq \lim_{\tau \rightarrow \infty} \left(\frac{n\tau^2}{n\tau^2+1}\right)^{d/2} \cdot C (n\tau^2+1)^{(d-5)/2} \cdot \left(1 / \int_{\Omega_0} e^{\frac{-\mu^T \Sigma^{-1} \mu}{2\tau^2}} d\mu\right)
\end{aligned}$$

As $\tau \rightarrow \infty$, $\int_{\Omega_0} e^{\frac{-\mu^T \Sigma^{-1} \mu}{2\tau^2}} d\mu$ approaches the volume of Ω_0 which is strictly positive and the first term $(\frac{n\tau^2}{n\tau^2+1})^{d/2}$ goes to 1. Moreover, since $d \leq 4$, the middle term $(n\tau^2 + 1)^{(d-5)/2}$ goes to zero. Consequently, we have:

$$\lim_{\tau \rightarrow \infty} \frac{\int_{\mathbb{R}^d} \int_{\mathbb{R}^d} \|\mu - \delta_{SA}^n(\xi)\|^2 p(\mu|\xi) m(\xi) d\mu d\xi - \int_{\mathbb{R}^d} \int_{\mathbb{R}^d} \|\mu - \delta_{Bayes}\|^2 p(\mu|\xi) m(\xi) d\mu d\xi}{\int_{\Omega_0} \pi_\tau(\mu) d\mu} = 0$$

By Theorem 5.1, δ_{SA}^n is admissible. \square

Proof of Theorem 1.3. The proof when X is a ball follows the same logic as the hypercube case of Theorem 1.2. Observe that now the Bayes' rule δ_{Bayes} differs from δ_{SA} only on the compact domain $\frac{n\tau^2+1}{n\tau^2} X$; outside this, both give the same point on the boundary of X . Thus, the numerator retains a factor of $(n\tau^2 + 1)^{-2}$, as opposed to the factor of $(n\tau^2 + 1)^{(d-5)/2}$ in the analysis of the hypercube case. Hence, for all $d \geq 1$, the ratio goes to 0 as $\tau \rightarrow \infty$. \square

6 Future Work

To the best of our knowledge, a thorough investigation of the admissibility of solution estimators for stochastic optimization problems has not been undertaken in the statistics or optimization literature. There are several avenues for continuing this line of investigation:

1. The most immediate question is whether the sample average solution for the quadratic objective subject to box constraints, as considered in this paper, continues to be admissible for dimension $d \geq 5$. We strongly suspect this to be true, but our current proof techniques are not able to resolve this either way. If one considers the James-Stein estimators for μ and then uses these to solve the constrained optimization problems, the standard arguments for inadmissibility break down because of the presence of the box constraints. The problem seems to be quite different, and significantly more complicated, compared to the unconstrained case that has been studied in classical statistics literature.
2. The next step, after resolving the higher dimension question, would be to consider general convex quadratic objectives $F(x, \xi) = x^T Q x + \xi^T x$ for some fixed positive (semi)definite matrix Q and the constraint X to be a general compact, convex set, as opposed to just box constraints. We believe new ideas beyond the techniques introduced in this paper are needed to analyze the admissibility of the sample average estimator for this convex quadratic program⁴. This problem is interesting from a financial engineering perspective, where the stochastic optimization problem seeks to minimize a coherent risk measure over a convex set. The simplest such measure is a weighted sum of the expectation and the variance of the returns, which can be modeled using the above $F(x, \xi)$.

⁴When Q is the identity and X is a scaled and translated ℓ_2 unit norm ball, as opposed to a box, the problem becomes equivalent to minimizing a linear function over the ball. Theorem 1.1 then applies to show admissibility in every dimension. This special case is also analyzed in [9].

3. One may also choose to avoid nonlinearities and stick to piecewise linear $F(x, \xi)$ and polyhedral X . Such objectives show up in the stochastic optimization literature under the name of *news-vendor type problems*. The current techniques of this paper do not easily apply directly to this setting either. In fact, in the simplest setting for the news-vendor problem, one has a function $F : \mathbb{R} \times \mathbb{R} \rightarrow \mathbb{R}$ given by $F(x, \xi) = cx - p \min\{x, \xi\} + r \max\{0, x - \xi\}$ and $X = [0, U]$ for known constants $0 < r < c < p$ and some given bound $U > 0$. In this setting, the natural distributions for ξ are not normal, but distributions whose support is contained in the nonnegative real axis. As a starting point, one can consider the uniform distribution setting where the mean of the uniform distribution, or the width of the uniform distribution or both are unknown.
4. For learning problems, such as neural network training with squared or logistic loss, what can be said about the admissibility of the sample average rule, which usually goes under the name of “empirical risk minimization”? Is the empirical risk minimization rule an admissible rule in the sense of statistical decision theory? It would be very interesting if the answer actually depends on the hypothesis class that is being learnt. It is also possible that decision rules that take the empirical risk objective and report a *local* optimum can be shown to dominate decision rules that report the global optimum, under certain conditions. This would be an interesting perspective on the debate whether local solutions are “better” in a theoretical sense than global optima.

Acknowledgments. We are extremely grateful to Prof. Daniel Naiman for helping us understand basic admissibility results from statistics and numerous follow-up discussions related to this work. He gave his time unquestioningly whenever we approached him. We also owe a great philosophical debt to Prof. Carey Priebe who always nudged us to think about optimization in statistical/probabilistic terms. His aphorism “The solution to an optimization problem is a random variable!” made us think about optimization in a new light. Prof. Priebe also had numerous discussions about the paper and gave unwavering encouragement and support. Prof. Jim Spall and Long Wang also gave useful input on a draft of the paper. Insightful comments and important pointers to prior work from two anonymous referees helped organize the results with more context. Finally, the main inspiration for this work came from listening to a beautiful talk by Prof. Gérard Cornuéjols explaining his work from [9] and [10].

References

- [1] James O Berger. *Statistical decision theory and Bayesian analysis*. Springer Science & Business Media, 2013.
- [2] Dimitris Bertsimas, Vishal Gupta, and Nathan Kallus. Data-driven robust optimization. *Mathematical Programming*, 167(2):235–292, 2018.
- [3] Peter J Bickel. Minimax estimation of the mean of a normal distribution when the parameter space is restricted. *The Annals of Statistics*, 9(6):1301–1309, 1981.

- [4] Peter J Bickel and Kjell A Doksum. *Mathematical statistics: basic ideas and selected topics, volume I*, volume 117. CRC Press, 2015.
- [5] John R Birge and Francois Louveaux. *Introduction to stochastic programming*. Springer Science & Business Media, 2011.
- [6] George Casella and William E Strawderman. Estimating a bounded normal mean. *The Annals of Statistics*, pages 870–878, 1981.
- [7] Alec Charras and Constance Van Eeden. Bayes and admissibility properties of estimators in truncated parameter spaces. *Canadian Journal of Statistics*, 19(2):121–134, 1991.
- [8] Leon Yang Chu, J George Shanthikumar, and Zuo-Jun Max Shen. Solving operational statistics via a bayesian analysis. *Operations Research Letters*, 36(1):110–116, 2008.
- [9] Danial Davarnia and Gérard Cornuéjols. From estimation to optimization via shrinkage. *Operations Research Letters*, 45(6):642–646, 2017.
- [10] Danial Davarnia, Burak Kocuk, and Gérard Cornuéjols. Bayesian solution estimators in stochastic optimization. http://www.optimization-online.org/DB_HTML/2017/11/6318.html, 2018.
- [11] Adam N Elmachtoub and Paul Grigas. Smart” predict, then optimize”. *arXiv preprint arXiv:1710.08005*, 2017.
- [12] Peyman Mohajerin Esfahani and Daniel Kuhn. Data-driven distributionally robust optimization using the wasserstein metric: Performance guarantees and tractable reformulations. *Mathematical Programming*, 171(1-2):115–166, 2018.
- [13] Dominique Fourdrinier and Éric Marchand. On bayes estimators with uniform priors on spheres and their comparative performance with maximum likelihood estimators for estimating bounded multivariate normal means. *Journal of Multivariate Analysis*, 101(6):1390–1399, 2010.
- [14] Dominique Fourdrinier, William E Strawderman, and Martin T Wells. *Shrinkage estimation*. Springer, 2018.
- [15] Constantine Gatsonis, Brenda MacGibbon, and William Strawderman. On the estimation of a restricted normal mean. *Statistics & probability letters*, 6(1):21–30, 1987.
- [16] John D Gorman and Alfred O Hero. Lower bounds for parametric estimation with constraints. *IEEE Transactions on Information Theory*, 36(6):1285–1301, 1990.
- [17] Vishal Gupta and Paat Rusmevichientong. Small-data, large-scale linear optimization with uncertain objectives. *Available at SSRN 3065655*, 2017.
- [18] John A Hartigan. Uniform priors on convex sets improve risk. *Statistics & probability letters*, 67(4):285–288, 2004.
- [19] Roger A. Horn and Charles R. Johnson. *Matrix Analysis*. John Wiley & Sons, second edition, 1985.

- [20] William James and Charles Stein. Estimation with quadratic loss. In *Proceedings of the fourth Berkeley symposium on mathematical statistics and probability*, volume 1, pages 361–379, 1961.
- [21] Iain Johnstone. On inadmissibility of some unbiased estimates of loss. *Statistical Decision Theory and Related Topics*, 4(1):361–379, 1988.
- [22] Ali Karimnezhad. Estimating a bounded normal mean relative to squared error loss function. *Journal of Sciences, Islamic Republic of Iran*, 22(3):267–276, 2011.
- [23] Somesh Kumar and Yogesh Mani Tripathi. Estimating a restricted normal mean. *Metrika*, 68(3):271–288, 2008.
- [24] Erich L Lehmann and George Casella. *Theory of point estimation*. Springer Science & Business Media, 2006.
- [25] Liwan H Liyanage and J George Shanthikumar. A practical inventory control policy using operational statistics. *Operations Research Letters*, 33(4):341–348, 2005.
- [26] Éric Marchand and François Perron. Improving on the mle of a bounded normal mean. *The Annals of Statistics*, 29(4):1078–1093, 2001.
- [27] Alexander Shapiro, Darinka Dentcheva, and Andrzej Ruszczyński. *Lectures on stochastic programming: modeling and theory*. SIAM, 2009.
- [28] Charles Stein. Inadmissibility of the usual estimator for the mean of a multivariate normal distribution. In *Proceedings of the third Berkeley symposium on mathematical statistics and probability*, volume 1, pages 197–206, 1956.
- [29] Aad W Van der Vaart. *Asymptotic statistics*, volume 3. Cambridge university press, 2000.
- [30] Bart PG Van Parys, Peyman Mohajerin Esfahani, and Daniel Kuhn. From data to decisions: Distributionally robust optimization is optimal. *arXiv preprint arXiv:1704.04118*, 2017.
- [31] Xianchao Xie, SC Kou, and Lawrence D Brown. Sure estimates for a heteroscedastic hierarchical model. *Journal of the American Statistical Association*, 107(500):1465–1479, 2012.